

Incentivizing Exploration with Heterogeneous Value of Money

Li Han

David Kempe

Ruixin Qiang

December 29, 2015

Abstract

Recently, Frazier et al. proposed a natural model for crowdsourced exploration of different a priori unknown options: a principal is interested in the long-term welfare of a population of agents who arrive one by one in a multi-armed bandit setting. However, each agent is myopic, so in order to incentivize him to explore options with better long-term prospects, the principal must offer the agent money. Frazier et al. showed that a simple class of policies called time-expanded are optimal in the worst case, and characterized their budget-reward tradeoff.

The previous work assumed that all agents are equally and uniformly susceptible to financial incentives. In reality, agents may have different utility for money. We therefore extend the model of Frazier et al. to allow agents that have heterogeneous and non-linear utilities for money. The principal is informed of the agent’s tradeoff via a signal that could be more or less informative.

Our main result is to show that a convex program can be used to derive a signal-dependent time-expanded policy which achieves the best possible Lagrangian reward in the worst case. The worst-case guarantee is matched by so-called “Diamonds in the Rough” instances; the proof that the guarantees match is based on showing that two different convex programs have the same optimal solution for these specific instances. These results also extend to the budgeted case as in Frazier et al. We also show that the optimal policy is monotone with respect to information, i.e., the approximation ratio of the optimal policy improves as the signals become more informative.

1 Introduction

The goal of mechanism design is to align incentives when different parties have conflicting interests. In the VCG mechanism, the mechanism designer wants to maximize social welfare whereas each bidder selfishly maximizes his own payoff. In revenue maximization, the objectives are even more directly opposed, as any increase in the bidders’ surplus hurts the revenue for the auctioneer. In all of these cases, it is the mechanism’s task to trade off between the differing interests.

The phrase “trade off” is also frequently applied in the context of online learning and the multi-armed bandit (MAB) problem, where the “exploration vs. exploitation tradeoff” is routinely referenced. However, in the traditional view of a single principal making a sequence of decisions to maximize long-term rewards, it is not clear what exactly is being traded off against what. Recent work by Frazier et al. [9] makes this tradeoff more explicit, by juxtaposing a principal (with a far-sighted goal of maximizing long-term rewards) with selfish and myopic agents. Thus, the principal wants to “explore,” while the agents want to “exploit.” In order to partially align the incentives, the principal can offer the agents monetary payments for pulling particular arms.

The framework of Frazier et al. [9] is motivated by many real-world applications, all sharing the property that the principal is interested in the long-term outcome of an exploration of different

options, but cannot carry out the exploration herself¹. Perhaps the most obvious fit is that of an online retailer with a large selection of similar products (e.g., cameras on amazon.com); in order to learn which of these products are best (and ensure that *future* buyers purchase the best product), the retailer needs to rely on customers to buy and review the products. Each customer prefers to purchase the best product for himself based on the current reviews, whereas the principal may want to obtain additional reviews for products that currently have few reviews, but may have the potential of being high quality. Customers can be incentivized to purchase such products by offering suitable discounts.

Other applications include crowd-sourced science projects, such as the search for celestial objects or bird or fish counts: individuals may prefer visiting areas with reliable sightings, while the principal would like underexplored areas to receive more coverage. In fact, even research funding can be naturally viewed in this context: while individual research groups may prefer to carry out research with good short-term rewards, funding agencies can use grants as an incentive to explore directions with long-term benefits.

Frazier et al. [9] explore this tradeoff under the standard time-discounted Bayesian² multi-armed bandit model (described formally in Section 2). In each round, each arm i has a known posterior reward distribution v_i conditioned on its history so far, and one arm is pulled based on the current state of the arms. The principal’s goal is to maximize the total expected time-discounted reward $R = \sum_{t=0}^{\infty} \gamma^t \mathbb{E}[v_{i_t}]$, where γ is the time discount factor. However, without incentives, each selfish agent would pull the *myopic* arm i maximizing the immediate expected reward $\mathbb{E}[v_i]$. When the principal offers payments c_i for pulling arms i , in [9], the agent’s utility for pulling arm i is $\mathbb{E}[v_i] + c_i$, and a myopic agent will choose the arm maximizing this sum.

Implicit in this model is the assumption that all agents have the same (one-to-one) tradeoff between arm rewards and payments. In reality, different agents might have different and non-linear tradeoffs between these two, due to a number of causes. The most obvious is that an agent with a large money endowment (personal wealth or research funding) may not value additional payments as highly as an agent with smaller endowment; this is generally the motivation for positing concave utility functions of money. In the case of an online retailer, another obvious reason is that different customers may intend to use the product for different amounts of time or with different intensity, making the optimization of quality more or less important. Concretely, a professional photographer may be much less willing to compromise on quality in return for a discount than an amateur.

The main contribution of the present article is an extension of the model and analysis of Frazier et al. [9] to incorporate non-uniform and non-linear tradeoffs between rewards and money. We assume that each myopic agent has a monotone and concave *money utility function* $\mu : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ mapping the agent’s payment to the corresponding utility. The utility of an agent with money utility function μ is quasi-linear: $\mathbb{E}[v_i] + \mu(c_i)$. The larger the values of μ , the easier it is to incentivize the agent with money, while an agent with $\mu \equiv 0$ cannot be incentivized at all. When an agent arrives in round t , we assume that his money utility function μ_t is drawn i.i.d. from some known distribution F .

An important question is then how much the principal knows about μ_t at the time she chooses the payment vector $\mathbf{c}_t = (c_{t,i})_i$ to announce for the arm pulls. In the worst case, the principal may know nothing about agent t as he arrives. In that case, the payment vector \mathbf{c}_t can only depend

¹To avoid ambiguity, we consistently refer to the principal as female and the agents as male.

²Both Frazier et al. [9] and our work in fact consider a generalization in which each arm constitutes an independent Markov chain with Martingale rewards.

on F . At the other extreme, the principal may learn the value of μ_t exactly. Then she is able to precisely control the myopic agent’s decision by setting \mathbf{c} accordingly³.

Reality will typically lie between these two extreme cases. Both financial endowments and intended use can be partially inferred from past searches and purchases in the case of an online retailer. This partial information will give the principal a more accurate estimate of the agent’s money utility function value than what could be learned from the prior distribution F alone, allowing her to better engineer the incentives.

We formally model the notion of partial information using the standard economic notion of an exogenous signaling scheme [30]. A signaling scheme specifies how signals are correlated with the ground truth, via a conditional distribution.

For each possible ground truth value (here: μ_t), the signaling scheme prescribes a distribution over possible signals $s \in \Sigma$ that could be revealed. Then, a known signaling scheme induces a posterior distribution when receiving a signal: when signal s is revealed to the principal, she can update her posterior belief of the money utility function μ from F to a more “accurate” $F(\mu|s)$.

The principal now has two goals, which stand in contrast with each other: minimizing her total expected time-discounted payment $C = \sum_{t=0}^{\infty} \gamma^t c_{t,it}$, and maximizing her total expected time-discounted reward R . These two quantities in some sense capture the two objectives that must be traded off: R is the long-term reward to be maximized, while C captures the loss in immediate payoffs. There are two natural ways of combining these two objectives: The first is to maximize a Lagrangian objective $R - \lambda C$, for some constant $\lambda \in (0, 1)$. The other is to maximize R subject to a constraint on the total expected time-discounted payments.

Our Results

In Section 7, we show that in a certain sense, *linear* functions μ constitute the worst case for the principal. Specifically, we show that for every distribution over functions μ and a corresponding signaling scheme, we can determine a distribution over linear functions such that the proposed mechanisms in this article perform at least as well under the original distribution as under the modified one. On the other hand, we show that for certain MAB instances, the original distribution does not allow strictly better mechanisms than the distribution over linear functions. The proof relies heavily on the techniques developed throughout the article; however, it retroactively justifies the sole focus in the rest of the article on the case that $\mu_t(x) = r_t \cdot x$.

Let OPT_γ be the optimal total expected time-discounted reward if the principal is allowed to pull arms herself. We call a policy an α -approximation policy if it achieves at least an α -fraction of OPT_γ for *every* MAB instance. (Precise definitions are given in Section 2.) The main result of this paper is to characterize the optimal way to utilize partial information of agents to incentivize exploration from them.

Theorem 1 *Let γ be the time discount factor. Given a prior distribution F (satisfying some technical conditions) and signaling scheme φ , one can efficiently compute a policy **TES** and $p^*(\varphi)$ such that the Lagrangian reward of **TES** is a $(1 - p^*(\varphi)\gamma)$ -approximation to OPT_γ . This bound is tight.*

Theorem 2 *Given a prior distribution F (again, satisfying some technical conditions), signaling scheme φ and budget constraint b , there exists a policy **TES** whose total expected time-discounted*

³If she wants arm i to be pulled, setting $c_i = \mu_t^{-1}(\max_j \mathbb{E}[v_j] - \mathbb{E}[v_i])$ suffices.

reward is a $\min_{\lambda}\{1 - p^(\varphi)\lambda + \lambda b\} - \epsilon$ approximation to OPT_{γ} , while spending at most $bOPT_{\gamma}$ in expectation. This bound is tight.*

In a sense, these theorems quantify the power of partial information φ about a money utility function distribution F in a single number $1 - p^*(\varphi)\gamma$, via the approximation guarantee that can be achieved using this signal. If this number is meaningful, more informative signaling schemes should allow for better approximation ratios. Specifically, a *garbling* [26] of a signaling scheme φ is another signaling scheme φ' whose output is computed solely from the *output* of φ , without knowledge of the true underlying state of the world. (A formal definition is given in Section 8.) In this sense, φ' cannot contain more information than φ . Then, we prove the following theorem in Section 8.

Theorem 3 *Let φ and φ' be two signaling schemes such that φ' is a garbling of φ . Then, $1 - p^*(\varphi)\gamma \geq 1 - p^*(\varphi')\gamma$.*

In Section 3, we prove the lower bound (algorithm) part of Theorem 1 by using the idea of time-expansion of policies from [9]. The main idea there is to randomize between pulls of the myopically optimal arm and the arm pulled by an optimal policy. Contrary to [9], we now have to carefully coordinate the randomization and payment policies for the different possible signals. This is accomplished by a convex program: the program is formulated predominantly as a heuristic, designed to cancel out myopic reward terms in the objective, which are otherwise difficult to analyze.

In Section 4, we prove that this heuristic is surprisingly optimal. For the matching upper bound, we use a class of instances called Diamonds-in-the-Rough [9], and show that the optimal policy using payments can only achieve a Lagrangian reward of $(1 - p^*(\phi)\gamma)OPT_{\gamma}$. We characterize the optimal policy with a different convex program. Using the Karush-Kuhn-Tucker (KKT) conditions for optimality of solutions to a convex program [24], we relate these two convex programs and show that their solution is actually the same, thus proving worst-case optimality.

We remark on the computational considerations of the proposed policies. Given an explicit representation of the distribution F and the signaling scheme, the convex program produces — without knowledge of any MAB instance — a vector of target randomization probabilities \mathbf{q} between myopic and non-myopic play, as well as the optimal p^* . When an actual MAB instance is specified, the algorithm draws on the optimum policy OPT_{η} for a different time-discount factor η . When an agent arrives, the principal needs to identify the myopically optimal arm and the arm pulled by OPT_{η} under a subset of the revealed information. Then, using the desired randomization, she can easily compute the required payments.

Thus, the only computationally challenging part is to compute the arm pulled by the optimum policy OPT_{η} . This can be accomplished by the well-known Gittins Index policy [11], which computes an index for each arm i based on the posterior distribution of rewards (or state of the Markov chain), and then chooses the arm with largest index. Thus, the policy avoids a combinatorial explosion in the number of arms, but computing the index of an arm can be non-trivial.

For the budgeted version, one first needs to find a suitable Lagrangian multiplier λ ensuring that the expected payment respects the budget, and then find the optimum policy for that λ . Finding λ requires analyzing the specific MAB instance and thus the optimal probabilities \mathbf{q} are not independent of the MAB instance any more. Theorem 2 is proved in Section 6.

Throughout this article and [9], it was assumed that the principal can observe the action that the agent took, and base the payment upon it. In many settings, the principal may only observe the arm *reward* that the agent obtained, but not the actual arm pulled. For instance, when the

agent is a scientist, a funding agency may be unable to tell whether a scientific result was achieved by deliberately following an ambitious agenda, or by stumbling upon it. In terms of the model, an inability to observe the action means that the payment scheme can be based only on the agent's observed reward. In Appendix B, we show that in this kind of scenario, it is impossible to incentivize the agent to pull optimal arms.

Related Work

The MAB problem was first proposed by Robbins [27] as a model for sequential experiments design. Under the Bayesian model with time-discounted rewards, the problem is solved optimally by the Gittins Index policy [11]; a further discussion is given in [32, 19, 10, 12].

An alternative objective, often pursued in the CS literature, is regret-minimization, as initiated by Lai and Robbins [22] within a Bayesian arm reward setting. Auer et al. [2, 3] gave an algorithm with regret bound for adversarial settings.

There is a rich literature that considers MAB problems when incentive issues arise. A common model is that a principal has to hire workers to pull arms, and both sides want to maximize their own utility. Singla and Krause [28] gave a truthful posted price mechanism. In [21, 25], the reward history is only known by the principal, and she can incentivize workers by disclosing limited information about the reward history to the worker. Ho et al. [18] used the MAB framework as a tool to design optimal contracts for agents with moral hazard. Using the technique of discretization, they achieved sublinear regret for the net utility (reward minus payment) over the time horizon. For a review of more work in the area, see the position paper [29].

Bergemann and Välimäki [6] and Bolton and Harris [7] consider incentive issues from another viewpoint. In their model, there are only two arms corresponding to two sellers and one or multiple buyers. Both sellers can set a price for a single pull of their own arm. After each pull, sellers can adjust their prices, and buyers can change their choices, which will eventually lead to an (efficient) equilibrium.

MAB problems with additional constraints and structure are also studied in various settings. Guha and Munagala [14, 16, 15] study a series of models with different constraints. In [20], for arms defined within a metric space and satisfying a Lipschitz condition, the authors found an optimal algorithm that matches the best possible regret ratio. Goel et al. [13] introduced an index policy, similar to the Gittins index policy, with constant approximation for the time horizon constrained MAB problem. Badanidiyuru et al. [4] investigated MAB problems with general multi-dimensional constraints.

The role of information in markets was introduced formally by Stigler [31]. Subsequently, Akerlof [1], Spence [30] began the study of effects of additional information, or signals, on the market. For exogenous signaling schemes, Hirshleifer [17], Bassan et al. [5], Lehrer et al. [23] explored the positive and negative effects on the equilibrium of different game settings.

2 Preliminaries

2.1 Multi-armed Bandits

In a Bayesian multi-armed bandits (MAB) instance, we are given N arms, each of which evolves independently as a known Markov chain whenever pulled. In each round⁴ $t = 0, 1, 2, \dots$, an algorithm can only pull one of the arms; the pulled arm will generate a random reward and then transition to a new state randomly according to the known Markov chain.

Formally, let $v_{t,i}$ be the random reward generated by arm i if it is pulled at time t . Let $S_{0,i}$ be the initial state of the Markov chain of the i -th arm and $S_{t,i}$ the state of arm i in round t . The distribution of $v_{t,i}$ is determined by $S_{t,i}$. Then, an MAB instance consists of N independent Markov chains and their initial states $\mathbf{S}_0 = (S_{0,i})_{i=1}^N$.

In this article, we are only interested in cases where the reward sequence for any single arm forms a *Martingale*, i.e.,

$$\mathbb{E}[\mathbb{E}[v_{t+1,i} \mid S_{t+1,i}] \mid S_{t,i}] = \mathbb{E}[v_{t,i} \mid S_{t,i}].$$

A *policy* \mathcal{A} is an algorithm that decides which arm to pull in round t based on the history of observations and the current state of all arms. Formally, a policy is a (randomized) mapping $\mathcal{A} : (t, \mathcal{H}_t, \mathbf{S}_t) \mapsto i_t$, where $\mathbf{S}_t = (S_{t,i})_{i=1}^N$ is the vector of arms' states, \mathcal{H}_t is the history up to time t , and i_t is the selected arm.

To evaluate the performance of a policy \mathcal{A} , we use standard time-discounting [11]. Let $\gamma \in (0, 1)$ be the time discount factor that measures the relative importance between future rewards and present rewards. If a policy \mathcal{A} receives an immediate reward of v_{t,i_t} in round t , it is discounted by a factor of γ^t and then added to the total reward. The total expected time-discounted reward can thus be defined as:

$$R^{(\gamma)}(\mathcal{A}) = \mathbb{E}_{\mathcal{A}} \left[\sum_{t=0}^{\infty} \gamma^t v_{t,i_t} \right],$$

where $\mathbb{E}_{\mathcal{A}}[\cdot]$ denotes the expectation conditioned on the policy \mathcal{A} being followed and the information it obtained, as in [9].

Given a time discount factor γ , we denote the optimal policy for that time discount (and also — in a slight overload of notation — its total expected time-discounted reward) by OPT_{γ} .

We call the arm with the maximum *immediate* expected reward $\mathbb{E}[v_{t,i} \mid \mathbf{S}_t]$ the *myopic arm*. A policy is called myopic if it pulls the myopic arm in each round. The myopic policy only *exploits* with no *exploration*, so it is inferior to the optimum policy in general, especially when the time-discount factor γ is close to 1.

2.2 Selfish Agents

We label each agent by the time t when he arrives. The Markov chain state \mathbf{S}_t and $\mathbb{E}[v_{t,i} \mid \mathbf{S}_t]$ are publicly known by both the agents and the principal.

In round t , the principal can offer a payment $c_{t,i}$ for pulling arm i . Incentivized by these extra payments $c_{t,i}$, agent t with money utility function μ_t now pulls the arm maximizing $\mathbb{E}[v_{t,i} \mid S_{t,i}] + \mu_t(c_{t,i})$. If the agent pulls arm i_t at time t , then the principal's reward from this pull is $\mathbb{E}[v_{t,i_t} \mid S_{t,i_t}]$,

⁴We use the terms “round” and “time” interchangeably.

and the agent's utility is $\mathbb{E}[v_{t,it} \mid S_{t,it}] + \mu_t(c_{t,it})$. ([9] studied the special case where $\mu_t(x) = x$ for all $x \geq 0$ and t .)

We assume a publicly known prior (whose distribution is denoted by F ⁵) over the money utility functions μ . When a new agent arrives, his money utility function is drawn from F independently of prior draws.

As discussed in the introduction, we show in Section 7 that the worst-case analysis can without loss of generality focus on the case in which all money utility function are linear, i.e., of the form $\mu_t(x) = r_t \cdot x$ for some r_t . Therefore, apart from that section itself, we will exclusively focus on the case of linear money utility functions. We then identify the distribution F with a distribution over the values r_t , which we call the *conversion ratio* of agent t . For the remainder of this article, all distributions and signals are assumed to be over conversion ratios instead of money utility functions.

Lemma 4 *For every distribution F over money utility functions, there exists another distribution F' over linear money utility functions (i.e., over conversion ratios) such that the optimal approximation ratio is the same for F and F' .*

The definition of *optimal* approximation ratio can be found in Definition 8.

2.3 Signaling Scheme

We assume the existence of an exogenous signaling scheme, i.e., the signaling scheme is given as input⁶. When an agent with conversion ratio r arrives, a signal $s \in \Sigma$ correlated to r is revealed to the principal according to the signaling scheme φ ; Σ is called the *signal space*, and we assume that it is countable. When the signal space is uncountable, defining the posterior probability density requires the use of Radon-Nikodym derivatives, and raises computational and representational issues. In Section 5, we consider what is perhaps the most interesting special case: that the signal reveals the precise value of r to the principal.

Formally, let $\varphi(r, s)$ be the probability that signal s is revealed when the agent's conversion ratio is r . In this way, the signals are statistically correlated with the conversion ratio r , and thus each signal reveals partial information about the true value of r . After receiving the signal s , the principal updates her posterior belief of the agent's conversion ratio according to Bayes Law⁷:

$$f_s(r) = \frac{\varphi(r, s)f(r)}{p_s}, \quad (1)$$

where $f_s(r)$ is the PDF of the posterior belief and $p_s = \int_0^\infty \varphi(r, s)f(r)dr$ is the probability that signal s is observed. For each signal $s \in \Sigma$, let F_s be the CDF of the corresponding posterior belief. As a special case, if the signaling scheme reveals no information, then $F_s = F$.

Throughout, we focus on the case when the posterior distributions F_s satisfies a condition called *semi-regularity* (this is the technical condition mentioned in Theorems 1 and 2), which is defined as follows:

Definition 5 (Semi-Regulariry) *A distribution with CDF G is called semi-regular if $\frac{1-x}{G^{-1}(x)}$ is convex. (When G is not invertible, we define $G^{-1}(x) := \sup\{t \geq 0 : G(t) \leq x\}$.)*

⁵When the money utility functions are always linear, F is also the cumulative distribution function of the slope.

⁶This is in contrast to the goal of *designing* a signaling scheme with certain properties.

⁷In Equation 1, if the support of r is finite, $f(r)$ can be replaced by the probability mass function.

Semi-regularity is a generalization of a well-known condition called regularity, defined as follows.

Definition 6 (Regularity) *A distribution with CDF G is regular if $G^{-1}(x) \cdot (1 - x)$ is concave.*

Lemma 7, proved in Appendix A, establishes that regularity implies semi-regularity, and hence our result is more general.

Lemma 7 *Let G be a CDF. If $G^{-1}(x) \cdot (1 - x)$ is concave, then $\frac{1-x}{G^{-1}(x)}$ is convex. In particular, regularity implies semi-regularity.*

2.4 Policies with partial information

The previous definition we gave of a policy did not take information about the agent's type into account. In light of this additional information, we give a refined definition. In addition to deciding which arm to pull, a policy may decide the payment to offer the agents based on the partial information obtained from signals. Formally, a policy is now a randomized mapping $\mathcal{A} : (t, \mathcal{H}_t, \mathbf{S}_t, s_t) \mapsto \mathbf{c}_t$, where s_t is the signal revealed in round t , and $c_{t,i}$ is the extra payment offered for pulling arm i in round t . After \mathbf{c}_t is announced, a myopic agent with conversion ratio r will pull the arm i_t that maximizes his own utility, causing that arm to transition according to the underlying Markov chain.

The expected payment of the principal is also time-discounted by the same⁸ factor γ . When \mathcal{A} is implemented, the total expected payment will be

$$C^{(\gamma)}(\mathcal{A}) = \mathbb{E}_{\mathcal{A}} \left[\sum_{t=0}^{\infty} \gamma^t c_{t,i_t} \right].$$

The principal faces two conflicting objectives: (a) maximizing the total expected time-discounted reward $R^{(\gamma)}(\mathcal{A})$; (b) minimizing the total expected time-discounted payment $C^{(\gamma)}(\mathcal{A})$; There are two natural ways of combining the two objectives: via a Lagrangian multiplier, or by optimizing one subject to a constraint on the other.

In the Lagrangian objective, the principal wishes to maximize $R^{(\gamma)}(\mathcal{A}) - \lambda C^{(\gamma)}(\mathcal{A})$ for some constant $\lambda \in (0, 1)$. Here, λ can also be regarded as the conversion ratio for the principal herself. Alternatively, the principal may be constrained by a budget b , and want to maximize $R^{(\gamma)}(\mathcal{A})$ subject to the constraint that $C^{(\gamma)}(\mathcal{A}) \leq b$.

2.5 Approximation Framework

Frazier et al. [9] performed a worst-case analysis over MAB instances and studied the (worst-case) approximation ratio. In this article, we similarly perform a worst-case analysis with respect to the MAB instances, while keeping an exogenous signaling scheme φ and the prior F fixed.

Definition 8 *For the Lagrangian objective of the problem, a policy \mathcal{A} has approximation ratio α under the signaling scheme φ and prior F if for all MAB instances⁹,*

$$R^{(\gamma)}(\mathcal{A}) - \lambda C^{(\gamma)}(\mathcal{A}) \geq \alpha \cdot \text{OPT}_{\gamma}. \quad (2)$$

⁸A natural justification for having the same discount factor is that after each round, with probability $1 - \gamma$, the game ends.

⁹Note that all $R^{(\gamma)}(\mathcal{A})$, $C^{(\gamma)}(\mathcal{A})$ and OPT_{γ} depend on the MAB instance.

We say that α is the optimal approximation ratio if there exists a policy with approximation ratio α and no policies have a better approximation ratio.

Likewise, for the budgeted version, a policy has approximation ratio α respecting budget b if

$$R^{(\gamma)}(\mathcal{A}) \geq \alpha \cdot OPT_\gamma \quad C^{(\gamma)}(\mathcal{A}) \leq b \cdot OPT_\gamma. \quad (3)$$

3 Lower bound: Time-Expanded Algorithm

In this section, we focus on the Lagrangian objective, and analyze *time-expanded algorithms*, in a generalization of the originally proposed notion of [9]. In a time-expanded algorithm, the principal randomizes between offering the agents no reward (having them play myopically), and offering the reward necessary to incentivize the agent to play the arm i_t^* according to a particular algorithm \mathcal{A} . In the presence of signals, the randomization probabilities for the different signals need to be chosen and optimized carefully, which is the main algorithmic contribution in this section. On the other hand, notice that if the posterior distribution of the conversion ratio conditioned on the signal is continuous, then the randomness in the user's type can instead be used as a randomization device, and the principal may be able to offer incentives deterministically.

More formally, Frazier et al. [9] define a time-expanded version $\mathbf{TE}_{p,\mathcal{A}}$ of a policy \mathcal{A} , parameterized by a probability p , as

$$\mathbf{TE}_{p,\mathcal{A}}(t) := \begin{cases} \mathcal{A}(\hat{\mathbf{S}}_t) & \text{if } Z_t = 1 \\ \arg\max_i \mathbb{E}[v_{t,i} \mid \mathbf{S}_t], & \text{otherwise} \end{cases}$$

where Z_t is a Bernoulli($1 - p$) variable. $\hat{\mathbf{S}}_t$ is the arm status that couples the execution of the time-expanded policy and the policy \mathcal{A} , which we will formally define later. When $Z_t = 1$, with the uniform agents defined in [9], in order to incentivize an agent to pull the non-myopic arm, the principal has to offer a payment of $\max_i \mathbb{E}[v_{t,i} \mid \mathbf{S}_t] - \mathbb{E}[v_{t,i_t^*} \mid \mathbf{S}_t]$, where $i_t^* = \mathcal{A}(\hat{\mathbf{S}}_t)$.

A time-expanded version of policy \mathcal{A} with signaling scheme φ works as follows: at time t , conditioned on the received signal s , the principal probabilistically offers a payment of c_{t,i_t^*} if the agent t pulls the arm i_t^* . Notice that only two options might maximize the agent's utility: pulling the myopic arm, or pulling the arm i_t^* and getting the payment. There is a direct correspondence between the payment c_{t,i_t^*} and the probability q_s that the agent chooses to pull the myopic arm. We will describe this correspondence below.

First, though, we discuss *which* arm i_t^* the principal is trying to incentivize the agent to pull. As in [9], it is necessary for the analysis that the execution of \mathcal{A} and of its time-expanded version can be coupled. To achieve this, in order to evaluate which arm should be pulled next by \mathcal{A} , the principal must only take the information obtained from the *non-myopic pulls* into consideration. Formally, we define $\hat{\mathbf{S}}_t$ as follows:

Define the random variable

$$Z_t := \begin{cases} 0 & \text{agent } t \text{ pulls the myopic arm} \\ 1 & \text{otherwise} \end{cases}$$

and $X_{t,i} = 1$ if arm i is pulled at time t and 0 otherwise. Notice that Z_t is a Bernoulli variable, and $\text{Prob}[Z_t = 0]$ depends on the received signal s and the payment offered by the principal. Let $N_{t,i} = \sum_{\tau=0}^{t-1} Z_\tau X_{\tau,i}$ be the number of non-myopic pulls of arm i before time t . Using this notation,

we define $\hat{S}_{t,i}$ to be the state of the Markov chain of arm i after the first¹⁰ $N_{t,i}$ pulls in the execution history of the time-expanded policy, and $\hat{\mathbf{S}}_t = (\hat{S}_{t,i})_i$.

Let F_s be the posterior CDF of the agent's conversion ratio. Let $x = \max_i \mathbb{E}[v_{t,i} \mid \mathbf{S}_t]$ be the expected reward of the myopic arm and $y = \mathbb{E}[v_{t,i_t^*} \mid \mathbf{S}_t]$ be the expected reward of arm i_t^* . If the principal offers a payment of c_{t,i_t^*} , then agents with conversion ratio $r < \frac{x-y}{c_{t,i_t^*}}$ will still choose the myopic arm. Assuming that agents break ties in favor of the principal, when $r \geq \frac{x-y}{c_{t,i_t^*}}$, they will prefer to pull arm i_t^* .

Conversely, in order to achieve a probability of q_s for pulling the myopic arm, the principal can choose a payment of $c_{t,i_t^*} = \inf\{c \mid F_s(\frac{x-y}{c}) \leq q_s\}$. If F_s is continuous at $\frac{x-y}{c_{t,i_t^*}}$, the probability of myopic play (conditioned on the signal) is exactly q_s , and c_{t,i_t^*} is the smallest payment achieving this probability. If there is a discontinuity at $\frac{x-y}{c_{t,i_t^*}}$, then for every $\epsilon > 0$, the probability of myopic play with payment $c_{t,i_t^*} + \epsilon$ is less than q_s . In that case, the principal offers a payment of c_{t,i_t^*} with probability $\frac{1-q_s}{1-F_s((x-y)/c_{t,i_t^*})}$ for pulling arm i_t^* , and no payment otherwise. Now, the probability of a myopic pull will again be exactly q_s .

To express the payment more concisely, we write $F_s^{-1}(q_s) = \sup\{r \mid F_s(r) \leq q_s\}$. Then, the payment can be expressed as $c_{t,i_t^*} = \frac{x-y}{F_s^{-1}(q_s)}$. In particular, when F_s is continuous, c_{t,i_t^*} will be offered deterministically; otherwise, the principal randomizes.

In summary, we have shown a one-to-one mapping between desired probabilities q_s for myopic play, and payments (and possibly probabilities, in the case of discontinuities) for achieving the q_s . We write $\mathbf{q} = (q_s)_{s \in \Sigma}$ for the vector of all probabilities. The unconditional (prior) probability of playing myopically is $\sum_{s \in \Sigma} p_s q_s$, and the expected payment $(x-y) \cdot \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)}$.

We can now summarize the argument above and give a formal definition of the time-expanded version of policy \mathcal{A} with signaling scheme φ ,

Definition 9 A policy $\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}$ is a time-expanded version of policy \mathcal{A} with signaling scheme φ , if at time t , after receiving the signal s about agent t 's conversion ratio, the principal chooses (randomized) payments such that the myopic arm is pulled with probability q_s , and the arm $i_t^* = \mathcal{A}(\hat{\mathbf{S}}_t)$ is pulled with probability $1 - q_s$.

This can be achieved by offering the agent, with probability $\frac{1-q_s}{1-\sup\{F_s(r) \mid F_s(r) \leq q_s\}}$, a payment of $\frac{\max_i \mathbb{E}[v_{t,i} \mid \mathbf{S}_t] - \mathbb{E}[v_{t,i_t^*} \mid \mathbf{S}_t]}{F_s^{-1}(q_s)}$ for pulling arm i_t^* .

Here, $F_s(r)$ is the CDF of the posterior distribution of the agent's conversion ratio conditioned on signal s .

The key technical lemma gives a sufficient condition on \mathbf{q} that allows us to obtain a good approximation ratio of the Lagrangian to the optimum solely in terms of $\sum_s p_s q_s$.

Lemma 10 Fix a signaling scheme φ . If \mathbf{q} satisfies $\sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)}$, there exists a policy \mathcal{A} , such that the time-expanded policy $\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}$ satisfies

$$R_\lambda^{(\gamma)}(\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}) \geq (1 - \gamma \cdot \sum_{s \in \Sigma} p_s q_s) \cdot \text{OPT}_\gamma. \quad (4)$$

¹⁰As in [9], in order to facilitate the analysis, this may include myopic and non-myopic pulls of arm i . For instance, if arm 1 was pulled as non-myopic arm at times 1 and 6, and a myopic pull of arm 1 occurred at time 3, then we would use the state of arm 1 after the pulls at times 1 and 3.

The proof of Lemma 10 rests heavily on variations of the following lemmas from [9]. They are extremely straightforward modifications of Lemmas 4.2 and 3.2 from [9], and the proofs are deferred to Appendix A.

Lemma 11 (Modification of Lemma 4.2 of [9]) *Given a parameter λ and a signaling scheme φ . Let $\zeta_{t-1} = \sum_{t' < t} Z_{t'}$ be the total number of non-myopic steps performed by the time-expanded algorithm $\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}$ prior to time t , where \mathbf{q} satisfies $\sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)}$. Then, for any $0 \leq n \leq t$,*

$$\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \zeta_{t-1} = n] \geq \mathbb{E}_{\mathcal{A}} [v_{n, i_n}].$$

Lemma 12 (Modification of Lemma 3.2 of [9]) *Given a parameter λ and a signaling scheme φ . Assume \mathbf{q} satisfies $\sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)}$, then for $\eta = \frac{(1-p)\gamma}{1-p\gamma}$, where $p = \sum_{s \in \Sigma} p_s q_s$, we have*

$$R_{\lambda}^{(\gamma)}(\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}) \geq \frac{1-\eta}{1-\gamma} \cdot R^{(\eta)}(\mathcal{A}).$$

Lemma 13 (Theorem 1.2 of [9]) *Consider a fixed MAB instance (without selfish agents) with two different time discount factors $\eta < \gamma$, and let $\text{OPT}_{\eta}, \text{OPT}_{\gamma}$ be the optimum time-discounted reward achievable under these discounts. Then, $\text{OPT}_{\eta} \geq \frac{(1-\gamma)^2}{(1-\eta)^2} \cdot \text{OPT}_{\gamma}$, and this bound is tight.*

Proof of Lemma 10. Let \mathcal{A} be the optimal policy for the time-discount factor $\eta = \frac{(1-p)\gamma}{1-p\gamma}$, where $p = \sum_{s \in \Sigma} p_s q_s$. We then have

$$\begin{aligned} R_{\lambda}^{(\gamma)}(\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}) &\stackrel{\text{Lemma 12}}{\geq} \frac{1-\eta}{1-\gamma} \cdot \text{OPT}_{\eta} \stackrel{\text{Lemma 13}}{\geq} \frac{1-\gamma}{1-\eta} \cdot \text{OPT}_{\gamma} \\ &= \frac{1-\gamma}{1-(1-p)\gamma/(1-p\gamma)} \cdot \text{OPT}_{\gamma} = (1-p\gamma) \cdot \text{OPT}_{\gamma}, \end{aligned}$$

completing the proof. ■

According to Lemma 10, the approximation guarantee of the time-expanded policy \mathbf{TES} is monotone decreasing in $p = \sum_{s \in \Sigma} p_s q_s$. This suggests a natural heuristic for choosing the myopic probabilities \mathbf{q} : minimize p subject to satisfying the conditions of the lemma. This optimization can be carried out using the following non-linear program. Surprisingly, this naïve heuristic, motivated predominantly by the need to cancel out terms in the proof of Lemma 10, actually gives us the optimal approximation ratio. We will prove this in Section 4.

$$\begin{aligned} &\text{minimize} && \sum_{s \in \Sigma} p_s q_s \\ &\text{subject to} && \sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)} \\ &&& 0 \leq q_s \leq 1, \quad \text{for all } s \in \Sigma. \end{aligned} \tag{5}$$

First, notice that the optimization problem is feasible, because $\mathbf{q} = \mathbf{1}$ is a trivial solution. Whenever F_s is semi-regular, $\frac{1-x}{F_s^{-1}(x)}$ is convex. Therefore, the feasibility region of the optimization problem (5) is convex, and the problem can be solved efficiently [8].

Theorem 14 *Given a signaling scheme φ , let \mathbf{q}^* be the optimal solution of the convex program (5), and p^* the optimal value. Let $\eta = \frac{(1-p^*)\gamma}{1-p^*\gamma}$. Then, $\mathbf{TES}_{\mathbf{q}^*, \text{OPT}_\eta}$ is a $(1 - p^*\gamma)$ -approximation policy to OPT_γ .*

This proves the first half of Theorem 1 in the introduction. Notice in Theorem 14 that \mathbf{q} can be determined without knowledge of the specific MAB instance; only the signaling scheme needs to be known.

4 Upper bound: Diamonds in the Rough

In this section, we show that the approximation ratio $1 - p^*\gamma$ is actually tight when the distribution F_s is semi-regular, where p^* is the value of the convex program (5). For simplicity, when \mathbf{q}^* is clear from the context, we let \mathbf{TES}^* denote the policy $\mathbf{TES}_{\mathbf{q}^*, \text{OPT}_\eta}$ where $\eta = \frac{(1-p^*)\gamma}{1-p^*\gamma}$ (as in Theorem 14). We will show that on a class of MAB instances called Diamonds-in-the-rough [9], the *optimal policy with payments* (defined below) can achieve only a $(1 - p^*\gamma)$ -fraction of OPT_γ . Therefore, not only is the analysis of \mathbf{TES}^* 's approximation ratio tight, but \mathbf{TES}^* also has the *optimal* approximation ratio $1 - p^*\gamma$.

Definition 15 *The Diamonds-in-the-rough MAB instance $\Delta(B, \gamma)$ is defined as follows. Arm 1 has constant value $1 - \gamma$. All other (essentially infinitely many) arms have the following reward distribution:*

1. *With probability $1/M$, the arm's reward is a degenerate distribution of the constant $(1-\gamma)B \cdot M$ (good state);*
2. *With probability $1 - 1/M$, the arm's reward is a degenerate distribution of the constant 0 (bad state).*

Note that if $B < 1$, then arm 1 is the myopic arm.

Since $\Delta(B, \gamma)$ is uniquely determined by B and is just one single instance, the optimal policy that maximizes the Lagrangian objective, i.e., $R^{(\gamma)}(\mathcal{A}) - \lambda C^{(\gamma)}(\mathcal{A})$, is well-defined¹¹. We call the policy that maximizes the Lagrangian objective the *optimal policy with payments*, and denote it by $\text{OPT}_\lambda^{(\gamma)}(\Delta(B, \gamma))$.

We can solve for the optimal policy with payments using another convex program, which we next derive. Suppose that the optimal policy with payments has time-discounted Lagrangian objective V . In the first round, it only has two options: (a) let the agent play myopically (i.e., pull the constant arm); (b) incentivize him to play a non-constant arm.

If option (a) is chosen and the agent pulled the constant arm, then the principal learns nothing and faces the same situation in the second round. So conditioned on the constant arm being pulled, the principal will get $1 - \gamma + \gamma V$. If option (b) is chosen and a non-constant arm was pulled, then with probability $1/M$, the non-constant arm will be revealed to be in the good state, and the principal does not need to pay any agent again, obtaining value $(1 - \gamma)B \cdot M \sum_{i=0}^{\infty} \gamma^i = B \cdot M$; with probability $1 - 1/M$, the non-constant arm will be revealed to be in the bad state, and the principal faces the same situation in the second round, obtaining value γV . Recall that c_s is the

¹¹This is in contrast to the case where the performance of a policy is evaluated on a *class* of instances rather than single instance.

payment needed to ensure that the myopic arm is played with probability at most q_s when signal s is revealed. To summarize, if we set the probabilities for myopic play to $(q_s)_{s \in \Sigma}$, then V satisfies the following equation:

$$V = (1 - \gamma + \gamma V) \sum_{s \in \Sigma} p_s q_s + \sum_{s \in \Sigma} p_s (1 - q_s) \left(\frac{1}{M} \cdot B \cdot M + \left(1 - \frac{1}{M}\right) \gamma V - \lambda c_s \right). \quad (6)$$

Solving for V while taking $M \rightarrow \infty$, we get $(1 - \gamma)V = (1 - \gamma) \sum_{s \in \Sigma} p_s q_s + \sum_{s \in \Sigma} p_s (1 - q_s) (B - \lambda c_s)$. As the difference between the expected rewards of the myopic arm and the non-myopic arm is $(1 - \gamma) - (1 - \gamma)B$, we have $c_s = \frac{(1 - \gamma)(1 - B)}{F_s^{-1}(q_s)}$. The optimal policy with payments needs to choose the best myopic probabilities, which is equivalent to:

$$\begin{aligned} & \text{maximize} && (1 - \gamma) \sum_{s \in \Sigma} p_s q_s + \sum_{s \in \Sigma} p_s (1 - q_s) \left(B - \lambda \frac{(1 - \gamma)(1 - B)}{F_s^{-1}(q_s)} \right) \\ & \text{subject to} && 0 \leq q_s \leq 1, \quad \text{for all } s \in \Sigma. \end{aligned} \quad (7)$$

Notice that the objective function of program (7) is concave, so the program is convex. Let $\hat{\mathbf{q}}$ be the optimal solution to the program (7). Denote by $\mathcal{A}(\hat{\mathbf{q}})$ the policy determined by $\hat{\mathbf{q}}$. Recall that \mathbf{q}^* is the solution to the following convex program:

$$\begin{aligned} & \text{minimize} && \sum_{s \in \Sigma} p_s q_s \\ & \text{subject to} && \sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1 - q_s}{F_s^{-1}(q_s)} \\ & && 0 \leq q_s \leq 1, \forall s \in \Sigma. \end{aligned} \quad (5)$$

Note that the $\hat{\mathbf{q}}$ are probabilities for choosing the myopic arm given by the above program and depend on a specific MAB instance, i.e., $\Delta(B, \gamma)$. On the other hand, \mathbf{q}^* is *independent* of any MAB instance and only depends on the signaling scheme φ and F . Lemma 18 shows that for the right choice of B , $\hat{\mathbf{q}}$ and \mathbf{q}^* actually coincide on the corresponding Diamonds-in-the-rough instance.

The proof relies heavily on the KKT condition [24], so we assume that $\frac{1-x}{F_s^{-1}(x)}$ is continuously differentiable and use the KKT condition for differentiable functions. When derivatives do not exist, we can use the sub-differential versions of KKT. Since we would like the characterization to hold for countably infinite signal spaces (and thus infinitely many variables), we have to be somewhat careful about the specific notion of differentiability, but note that standard notions such as Gateaux Differentiability [24] can be used here.

Definition 16 (Slater condition) *For a convex program with no equality constraints, the Slater condition is met if there exists a point \mathbf{x} such that $g_i(\mathbf{x}) < 0$ for all constraints i .*

It is easy to check that the program (5) satisfies the Slater condition.

Theorem 17 (KKT condition for infinite dimension) *Consider a program satisfying the Slater condition: minimize $f(\mathbf{x})$ subject to $g_i(\mathbf{x}) \leq 0$ for $i = 1, \dots, m$. If \mathbf{x}^* is the local minimum, then*

there exist multipliers μ_i for $i = 1, \dots, m$ such that:

$$\begin{aligned} \nabla f(\mathbf{x}^*) + \sum_{i=1}^m \mu_i \nabla g_i(\mathbf{x}^*) &= 0 \\ g_i(\mathbf{x}^*) &\leq 0 \\ \mu_i &\geq 0 \\ \mu_i \cdot g_i(\mathbf{x}^*) &= 0 \end{aligned} \tag{8}$$

Lemma 18 *There exists a B such that the myopic probabilities given by the convex program (5) are equal to the myopic probabilities given by program (7).*

Proof. First, we observe that the boundary case $q_s = 0$ can be safely ignored in both convex programs (5) and (7). This is because $\frac{1}{F_s^{-1}(q)}$ approaches infinity as $q_s \rightarrow 0$. For program (5), this violates the non-trivial feasibility constraint; for program (7), this is clearly sub-optimal.

Next, we prove that every local optimum \mathbf{q}^* is interior for convex program (5). Applying the stationarity condition of the KKT Theorem to program (5), we know that for every $s \in \Sigma$, the optimal solution \mathbf{q}^* satisfies:

$$-p_s = \mu_s + \sigma \left(\lambda p_s \frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} \Big|_{x=q_s^*} - p_s \right). \tag{9}$$

Here, μ_s is the multiplier for the constraint $0 \leq q_s \leq 1$, and σ is the multiplier for the non-trivial constraint. It is important that σ is a constant that is independent of any MAB instance.

Setting $q_s^* = 1$ for all $s \in \Sigma$ is clearly not an optimal solution as there will be slack in the constraints and decreasing q_s^* can improve the objective for program (5). Hence, there is an $s \in \Sigma$ such that $q_s^* < 1$. By complementarity, $\mu_s = 0$, so that $-p_s = \sigma \left(\lambda p_s \frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} \Big|_{x=q_s^*} - p_s \right)$. Canceling p_s , we have $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} \Big|_{x=q_s^*} = \frac{\sigma-1}{\lambda\sigma}$. But we know that F_s is semi-regular, so $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} \Big|_{x=q_s^*} < \frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} \Big|_{x=1} = 0$. (Equality is impossible, as it would imply $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} = 0$ for $x \in [q_s^*, 1]$ by definition of semi-regularity. But then, $\frac{1-x}{F_s^{-1}(x)}$ would not be strictly monotone, violating the fact that F_s is non-decreasing.) Because $\frac{\sigma-1}{\lambda\sigma} < 0$, we infer that $\sigma \in (0, 1)$.

On the other hand, if $q_s^* = 1$ for some $s \in \Sigma$, then $\mu_s \geq 0$ and $-p_s = \mu_s - \sigma \cdot p_s$ (as $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} \Big|_{x=1} = 0$), so $\sigma \geq 1$. This would be a contradiction to $\sigma < 1$, so we infer that $q_s^* < 1$ for every $s \in \Sigma$.

For program (7), we first compute the partial derivative of the objective w.r.t. q_s ; it is

$$p_s \left(1 - \gamma - \lambda(1 - \gamma)(1 - B)(1 - q_s) \frac{\partial}{\partial x} \frac{1}{F_s^{-1}(x)} \Big|_{x=q_s} - B + \lambda \frac{(1 - \gamma)(1 - B)}{F_s^{-1}(x)} \right). \tag{10}$$

Rearranging and using the fact that $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} = (1-x) \frac{\partial}{\partial x} \frac{1}{F_s^{-1}(x)} - \frac{1}{F_s^{-1}(x)}$, we can rewrite it as

$$p_s \left(1 - \gamma - \lambda(1 - \gamma)(1 - B) \frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)} \Big|_{x=q_s} - B \right). \tag{11}$$

Let $B \in (1 - \gamma, 1)$ solve the equation $\frac{1-\gamma-B}{(1-\gamma)(1-B)} = \frac{\sigma-1}{\lambda\sigma}$. A solution exists because $\sigma \in (0, 1)$ and $\frac{1-\gamma-B}{(1-\gamma)(1-B)}$ varies continuously from 0 to $-\infty$ as B varies from $1 - \gamma$ to 1. Notice that the constant arm is still the myopic choice when $B < 1$.

Recall that $\hat{\mathbf{q}}$ is a global, and thus local, maximum of program (7). Assume that $\hat{q}_s = 1$ for some s . Then, $\frac{\partial}{\partial q_s} \frac{1-x}{F_s^{-1}(x)}|_{x=\hat{q}_s} = 0$, so the partial derivative of the objective (11) w.r.t. q_s will be $p_s(1 - \gamma - B)$, which is negative because $B > 1 - \gamma$. Thus, decreasing \hat{q}_s would maintain feasibility and increase the value of the program (7). Hence, any local maximum of the program (7) must be an interior point as well.

In summary, both \mathbf{q}^* and $\hat{\mathbf{q}}$ are interior. By the KKT condition for \mathbf{q}^* , all the μ_i 's are zero, and thus $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)}|_{x=q_s^*} = \frac{\sigma-1}{\lambda\sigma}$. Because $\hat{\mathbf{q}}$ is a local optimum, expression (11) must be equal to zero, so $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)}|_{x=\hat{q}_s} = \frac{1-\gamma-B}{(1-\gamma)(1-B)}$.

By the choice of B , we then obtain that $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)}|_{x=q_s^*} = \frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)}|_{x=\hat{q}_s} = \frac{1-\gamma-B}{(1-\gamma)(1-B)}$. By semi-regularity of F_s , $\frac{\partial}{\partial x} \frac{1-x}{F_s^{-1}(x)}$ is negative and non-decreasing. Thus, it must be constant on the entire interval $[\hat{q}_s, q_s^*]$.

This means that the derivative (11) of the objective of program (7) is constant over the entire interval, and we established above that it is 0 at \hat{q}_s . Thus, the objective function of program (7) is unchanged by replacing \hat{q}_s with q_s^* . By performing this operation for all s , we eventually obtain that without loss of generality, $\hat{\mathbf{q}} = \mathbf{q}^*$. ■

Based on this lemma, we now prove the main theorem in this section. This also proves the second half of Theorem 1 in the introduction.

Theorem 19 *The policy \mathbf{TES}^* , parameterized by \mathbf{q}^* , has optimal approximation ratio $1 - p^*\gamma$. In particular, there exists a worst-case MAB instance in which the optimal policy with payments achieves exactly a Lagrangian reward of a $(1 - p^*\gamma)$ fraction of the optimum.*

Proof. Lemma 18 showed that for every signaling scheme, there is a B (a parameter of the non-constant arms) such that \mathbf{q}^* and $\hat{\mathbf{q}}$ coincide on the MAB instance $\Delta(B, \gamma)$. It remains to show that these instances are in fact worst-case instances, i.e., that the ratio between $R_\lambda^{(\gamma)}(\mathcal{A}(\hat{\mathbf{q}}))$ and OPT_γ in $\Delta(B, \gamma)$ is exactly $1 - p^*\gamma$.

First note that $\text{OPT}_\gamma = \frac{B}{1-\gamma}$: once a non-constant arm is revealed to be in a good state, it is optimal to pull that arm forever. We next show that $R_\lambda^{(\gamma)}(\mathcal{A}(\hat{\mathbf{q}})) = \frac{B(1-p^*\gamma)}{1-\gamma}$. Thereto, we use both convex programs and the fact that $\hat{\mathbf{q}} = \mathbf{q}^*$. Substituting \mathbf{q}^* into the objective of program (7), we obtain the following Lagrangian reward for the time-expanded policy:

$$R_\lambda^{(\gamma)}(\mathcal{A}(\hat{\mathbf{q}})) = \frac{1}{1-\gamma} \left((1-\gamma) \sum_{s \in \Sigma} p_s q_s^* + \sum_{s \in \Sigma} p_s (1 - q_s^*) \left(B - \lambda \frac{(1-\gamma)(1-B)}{F_s^{-1}(q_s^*)} \right) \right).$$

The non-trivial constraint of program (5) has to be tight; otherwise, the objective could be increased by lowering individual q_s^* values. Therefore, $\sum_{s \in \Sigma} p_s q_s^* = \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s^*}{F_s^{-1}(q_s^*)}$, which allows us to simplify the previous expression to

$$R_\lambda^{(\gamma)}(\mathcal{A}(\hat{\mathbf{q}})) = \frac{B(1-p^*\gamma)}{1-\gamma} = (1-p^*\gamma) \cdot \text{OPT}_\gamma.$$

Thus, on this instance, $R_\lambda^{(\gamma)}(\mathbf{TES}^*) \leq (1 - p^*\gamma)\text{OPT}_\gamma$. By Theorem 14, this bound is tight. ■

5 Full Information Revelation

Our main positive results hold for the case of countable or finite signal spaces, whereas uncountable signal spaces lead to technical challenges. However, one important special case of uncountable signal spaces is more easily handled, namely, when the principal learns the exact conversion ratio r , i.e., $s = r$. We show that in that case, r itself can be used as the sole randomization device, leading to a *threshold policy*. In this section, we assume that the distribution F is continuous (an assumption that was not needed in Section 3).

5.1 Optimal Time-Expanded Policy

Our first goal will be to show that the optimal time-expanded policy fixes a threshold θ and only incentivizes agents whose conversion ratio lies above the threshold. Then, an optimization over threshold policies is easy to carry out.

Definition 20 *The threshold policy $\mathbf{TP}_{\theta, \mathcal{A}}$ with threshold θ is defined as follows: When an agent with conversion ratio r arrives, he is incentivized with suitable payment to pull $i_t^* = \mathcal{A}(\hat{\mathbf{S}}_t)$ if and only if $r \geq \theta$.*

Lemma 21 *Consider a single arm pull, and a value $q \in [0, 1]$. Among all policies that have this arm pull be myopic with probability q , the one minimizing expected cost is a threshold policy.*

Proof. Consider any policy \mathcal{P} , and assume that $q = \text{Prob}[\mathcal{P} \text{ lets the agent play myopically}]$. Now, define $\theta = \sup\{r | F(r) = q\}$, so that $F(\theta) = q$.

As before, we write $x = \max_i \mathbb{E}[v_{t,i} | \mathbf{S}_t, Z_{0:t-1}]$ and $y = \mathbb{E}[v_{t,i_t^*} | \mathbf{S}_t, Z_{0:t-1}]$.

For any r , let $P_r = 1$ iff \mathcal{P} incentivizes agents with conversion ratio r . The expected payment of policy \mathcal{P} is then

$$\begin{aligned} (x - y) \cdot \int_0^\infty \frac{P_r}{r} dF(r) &= (x - y) \cdot \left(\int_\theta^\infty \frac{1}{r} dF(r) + \int_0^\theta \frac{P_r}{r} dF(r) - \int_\theta^\infty \frac{1 - P_r}{r} dF(r) \right) \\ &\geq (x - y) \cdot \left(\int_\theta^\infty \frac{1}{r} dF(r) + \int_0^\theta \frac{P_r}{\theta} dF(r) - \int_\theta^\infty \frac{1 - P_r}{\theta} dF(r) \right) \\ &= (x - y) \cdot \int_\theta^\infty \frac{1}{r} dF(r), \end{aligned}$$

where the last step used the definition of θ , implying that the measure (under F) of conversion ratio $r < \theta$ for which \mathcal{P} incentivizes agents is the same as the measure of $r > \theta$ for which \mathcal{P} does not incentivize agents.

The final expression is the expected Lagrangian cost of \mathcal{P}' , so we have shown that \mathcal{P}' has no larger cost than \mathcal{P} in this one round. \blacksquare

Lemma 22 *The Lagrangian objective of any time-expanded policy \mathcal{P} of \mathcal{A} is (weakly) dominated by that of a threshold policy.*

Proof. Consider a time-expanded policy \mathcal{P} . Thus, the probability q with which \mathcal{P} lets the agent play myopically is the same in each round. As in the proof of Lemma 21, we let $\theta = \sup\{r | F(r) = q\}$.

Because $F(\theta) = q$, in each round, the reward of \mathcal{P} is the same as that of the threshold policy with threshold θ .

Lemma 21 establishes that in each round, the cost of \mathcal{P}' is no more than that of \mathcal{P} ; thus, the Lagrangian objective for \mathcal{P}' is at least as large as for \mathcal{P} . \blacksquare

Because of Lemma 22, it suffices to study the optimal threshold policy, and determine the correct threshold.

As before, we let $x = \max_i \mathbb{E}[v_{t,i} \mid \mathbf{S}_t, Z_{0:t-1}]$, and $y = \mathbb{E}[v_{t,i_t^*} \mid \mathbf{S}_t, Z_{0:t-1}]$. The expected Lagrangian reward of $\mathbf{TP}_{\theta, \mathcal{A}}$ is

$$x \cdot F(\theta) + \int_{\theta}^{\infty} y - \frac{\lambda(x-y)}{r} dF(r) = y + (x-y) \cdot (F(\theta) - \lambda \int_{\theta}^{\infty} \frac{dF(r)}{r}).$$

Because F was assumed continuous, so is $H(\theta) = F(\theta) - \lambda \int_{\theta}^{\infty} \frac{dF(r)}{r}$; and because $H(0) < 0$, $\lim_{\theta \rightarrow \infty} H(\theta) = 1$, the equation $F(\theta) = \lambda \int_{\theta}^{\infty} \frac{dF(r)}{r}$ has a solution θ . Fixing this choice of θ , the expected Lagrangian payoff simplifies to y . The rest of the proof now proceeds as in the proof of Lemma 10, yielding an approximation ratio of $1 - F(\theta)\gamma$. Writing $p^* = F(\theta)$ for the unconditional probability of a myopic pull under $\mathbf{TP}_{\theta, \mathcal{A}}$, we obtain the same $(1 - p^*\gamma)$ approximation ratio for the Lagrangian objective as for the case of discrete signals.

5.2 Upper Bound

As for discrete signals, we next give a Diamonds-in-the-rough instance $\Delta(B, \gamma)$ on which the upper bound for any policy matches the approximation ratio of the threshold policy. Consider the choice of the policy in the first round; it allows the agent to play myopically with some probability q . By Lemma 21, the optimal way to implement this probability q is to choose a threshold¹² θ and offer incentives to the agent if and only if $r \geq \theta$. Denote by θ^* the threshold we get for the time-expanded policy, i.e., the solution of $F(\theta) = \lambda \int_{\theta}^{\infty} \frac{dF(r)}{r}$. Now, similar to Equation (6), the corresponding Lagrangian objective is

$$V = (1 - \gamma + \gamma V)F(\theta) + \int_{\theta}^{\infty} \left(\frac{1}{M} \cdot B \cdot M + (1 - \frac{1}{M})\gamma V - \lambda \frac{(1 - \gamma)(1 - B)}{r} \right) dF(r). \quad (12)$$

Letting $M \rightarrow \infty$ and solving Equation (12), we obtain

$$(1 - \gamma)V = (1 - \gamma - B)F(\theta) + B - \lambda(1 - \gamma)(1 - B) \int_{\theta}^{\infty} \frac{dF(r)}{r}. \quad (13)$$

Taking a derivative of Equation (13) with respect to θ suggests (note that the function may not be differentiable, so this is merely used as a tool to suggest a useful choice) that if we set B to solve $-\lambda(1 - \gamma)(1 - B) = (1 - \gamma - B)\theta^*$, then our previously chosen threshold θ^* will be optimal for the instance $\Delta(B, \gamma)$.

We next verify that this is indeed the case. Substituting the value of B into Equation (13), we obtain that maximizing Equation (13) is equivalent to minimizing $G(\theta) = \frac{1}{\theta^*}F(\theta) + \int_{\theta}^{\infty} \frac{dF(r)}{r}$. Now, for any θ , we have

¹²Note that a priori, it is not clear that this threshold will not change in subsequent rounds; hence, we cannot yet state that a threshold policy is optimal.

$$\begin{aligned}
G(\theta^*) - G(\theta) &= \left(\frac{F(\theta^*)}{\theta^*} + \int_{\theta^*}^{\infty} \frac{dF(r)}{r} \right) - \left(\frac{F(\theta)}{\theta^*} + \int_{\theta}^{\infty} \frac{dF(r)}{r} \right) \\
&= \frac{F(\theta^*) - F(\theta)}{\theta^*} + \int_{\theta^*}^{\theta} \frac{dF(r)}{r} \\
&\leq \frac{F(\theta^*) - F(\theta)}{\theta^*} + \frac{F(\theta) - F(\theta^*)}{\theta^*} = 0.
\end{aligned}$$

So θ^* is the maximizer of Equation (13) for the specific choice of B .

Thus, on this particular instance, the ratio achieved by our threshold policy matches that of best possible policy.

6 Budgeted version

In this section, we show matching lower and upper bounds for maximizing the total expected time-discounted reward with a budget constraint. Let b be the fraction of OPT_{γ} that the principal is allowed to use (i.e., the principal's budget is $b \cdot OPT_{\gamma}$). Recall that $p^*(\lambda)$ is the optimal value of the convex program (5).

Theorem 23 *Given budget $b \cdot OPT_{\gamma}$, there exists a policy whose approximation ratio (with respect to OPT_{γ}) is*

$$\min_{\lambda} \{1 - p^*(\lambda)\gamma + \lambda b\}.$$

Proof. First, one can prove that for every $p \in [0, 1]$, there exists a λ such that $p^*(\lambda) = p$. Then, as every unconditional myopic probability p can be achieved using some λ , we can follow the same approach in Frazier et al. [9] by taking the limit of a sequence of probabilities $(p_n)_n$ that approaches p . Here, the policy corresponding to each p_n respects (or exhausts) the budget, while the policy for p exhausts (or respects) the budget. By suitable randomization between these policies, one can show that the approximation ratio (with respect to OPT_{γ}) approaches $\min_{\lambda} \{1 - p^*(\lambda) + \lambda b\}$ in the limit. ■

Theorem 24 *Given budget $b \cdot OPT_{\gamma}$, the factor $\min_{\lambda} \{1 - p^*(\lambda)\gamma + \lambda b\}$ is tight.*

Proof. Again consider the class of Diamonds-in-the-rough instances. Assume that the optimal policy that respects the budget is \mathcal{A} . Thus $C^{(\gamma)}(\mathcal{A}) \leq b \cdot OPT_{\gamma}$ for every MAB instance.

Moreover, by Lemma 18, we know that for every Lagrangian multiplier λ , the optimal policy for that λ and its corresponding $\Delta(B, \gamma)$ has Lagrangian objective value exactly $(1 - p^*(\lambda)\gamma) \cdot OPT_{\gamma}$.

This means that for every Lagrangian multiplier λ (and its corresponding Diamonds-in-the-rough instance $\Delta(B, \gamma)$), we have the following equivalent inequalities for all $\lambda \in (0, 1)$,

$$\begin{aligned}
R^{(\gamma)}(\mathcal{A}) - \lambda C^{(\gamma)}(\mathcal{A}) &\leq (1 - p^*(\lambda)\gamma) \cdot OPT_{\gamma} \\
R^{(\gamma)}(\mathcal{A}) &\leq (1 - p^*(\lambda)\gamma) \cdot OPT_{\gamma} + \lambda \cdot C^{(\gamma)}(\mathcal{A}) \\
R^{(\gamma)}(\mathcal{A}) &\leq (1 - p^*(\lambda)\gamma + \lambda b) \cdot OPT_{\gamma}
\end{aligned}$$

So

$$R^{(\gamma)}(\mathcal{A}) \leq \min_{\lambda} \{(1 - p^*(\lambda)\gamma + \lambda b)\} \cdot OPT_{\gamma}.$$

This matches the theorem above. ■

7 Reducing Concave Money Value to Linear Money Value

In this section, we show that among concave money utility functions, linear functions obtain the worst-case approximation ratio. This shows that our assumption that agents have linear money utility functions is without loss of generality in regard to worst-case approximation ratios. We fix the distribution (or prior) over agents and the signaling scheme and then perform a worst-case analysis over MAB instances. The intuition is that the rewards of an MAB instance can be scaled up so large that only the asymptotic behavior of concave money utility functions matters.

For any money utility function μ , define $r_\mu = \lim_{x \rightarrow \infty} \frac{\mu(x)}{x}$ to be the limiting slope¹³ of μ , and $d_\mu = \inf\{t \geq 0 \mid r_\mu x + t \geq \mu(x) \text{ for all } x \geq 0\}$ the minimum intercept. Thus, we ensure that for all x :

$$r_\mu x \leq \mu(x) \leq d_\mu + r_\mu x.$$

Fix a signaling scheme φ . For a given signal s and its original posterior distribution F_s over concave functions, define two distributions over affine functions. Under \mathcal{D}_s^1 , each function μ is mapped to the linear function $x \mapsto r_\mu \cdot x$. Under \mathcal{D}_s^2 , the function μ is mapped to the function $x \mapsto d_\mu + r_\mu \cdot x$. We assume that \mathcal{D}_s^1 is a continuous distribution for each s .

Our proof breaks into two parts: lower bound (algorithm) and upper bound (impossibility).

Lower bound: Given the true distribution $(F_s)_{s \in \Sigma}$ over concave functions, instead compute the optimal policy for $(\mathcal{D}_s^1)_{s \in \Sigma}$, which we denote by OPT_1 . Notice that OPT_1 is a time-expansion policy and only offers payment on the Gittins arm in each round.

Because $r_\mu \cdot x \leq \mu(x)$ for all x , each agent likes money more under μ than under r_μ . Thus, we can couple the choices of agents between the two distributions: each agent who plays non-myopically under the policy for \mathcal{D}^1 can also be incentivized to do so under F using lower payments in expectation. Therefore, the optimal policy for the true distribution can only do better.

Formally, let arm i^* be the Gittins arm at some round t during the execution of OPT_1 . Suppose that OPT_1 offers a payment of p on arm i^* and the arm will be pulled with probability q when agents are drawn from \mathcal{D}_s^1 . When agents are drawn from F_s instead, the probability q_F that the arm will be pulled is no less than q as agents like money more under μ . Then our new policy for F_s can randomize between offering the same payment p or offering no payment so that the overall probability of pulling i^* under F_s is the same as that under \mathcal{D}_s^1 . In this way, we can couple the execution of OPT_1 and our new policy so that agents always make a same choices yet OPT_1 offers more payment in expectation.

Upper bound: Next we show that in the worst case, the distribution (F_s) over concave functions does not yield a better approximation guarantee than (\mathcal{D}_s^1) . First, notice that by the same argument as the previous paragraph, the principal's utility under (F_s) is upper bounded by her utility with the distribution $(\mathcal{D}_s^2)_{s \in \Sigma}$ over affine functions. We will upper-bound the worst-case utility under (\mathcal{D}_s^2) in terms of that under (\mathcal{D}_s^1) .

Thereto, we modify the Diamonds-in-the-rough instance from Section 4, by scaling all reward values by some large constant C . Let $\xi = C(1 - \gamma)(1 - B)$ be the difference between the reward of the myopic arm and the Gittins arm. We will prove that for some carefully chosen constant C , the optimal policy with payments for \mathcal{D}^2 cannot achieve much more than the optimum for \mathcal{D}^1 . Let $(c_s)_s$ be the payments offered under different signals under the optimal solution for \mathcal{D}^2 .

¹³As μ is concave, this limit always exists and is finite.

By the argument preceding the program (7), the optimum solution for \mathcal{D}^2 maximizes

$$C(1-\gamma) \cdot \sum_s p_s \cdot \text{Prob}_{(r,d) \sim \mathcal{D}_s^2}[r \cdot c_s + d \leq \xi] + \sum_s p_s \cdot (1 - \text{Prob}_{(r,d) \sim \mathcal{D}_s^2}[r \cdot c_s + d \leq \xi])(C \cdot B - \lambda c_s), \quad (14)$$

while the optimum solution for \mathcal{D}^1 maximizes

$$C(1-\gamma) \cdot \sum_s p_s \cdot \text{Prob}_{r \sim \mathcal{D}_s^1}[r \cdot c_s \leq \xi] + \sum_s p_s \cdot (1 - \text{Prob}_{r \sim \mathcal{D}_s^1}[r \cdot c_s \leq \xi])(C \cdot B - \lambda c_s). \quad (15)$$

Now, for any arbitrarily small ϵ , focus on a finite set Σ' of signals which together contribute at least a $(1 - \epsilon)$ fraction of the optimum utility under both \mathcal{D}^1 and \mathcal{D}^2 . Note that such a set must always exist, because we assumed that the signal space is countable; by sorting the signals by their contribution and taking prefixes in this ordering, we see that the total utility of these prefixes must converge to the overall total utility as the size of the prefix grows.

For each individual signal s , the optimal $c_s \rightarrow \infty$ as $C \rightarrow \infty$. Thus, because Σ' is finite, for any desired lower bound α , there exists a C such that $c_s \geq \alpha$ for all $s \in \Sigma'$. Similarly, for any $\epsilon > 0$, there is a (sufficiently large) β such that $\text{Prob}_{(r,d) \sim \mathcal{D}_s^2}[d \leq \beta] \geq 1 - \epsilon$ holds simultaneously for all $s \in \Sigma'$. Given a target ϵ , we first choose a suitable β , and then choose C to ensure that $\alpha \gg \beta$. Then, we get that for all $s \in \Sigma'$,

$$\begin{aligned} \text{Prob}_{(r,d) \sim \mathcal{D}_s^2}[rc_s + d \leq \xi] &= \text{Prob}_{(r,d) \sim \mathcal{D}_s^2}\left[r \leq \frac{\xi - d}{c_s}\right] \\ &\geq (1 - \epsilon) \cdot \text{Prob}_{(r,d) \sim \mathcal{D}_s^2}\left[r \leq \frac{\xi}{c_s} - \frac{\beta}{\alpha}\right] \\ &= (1 - \epsilon) \cdot \text{Prob}_{r \sim \mathcal{D}_s^1}\left[r \leq \frac{\xi}{c_s} - \frac{\beta}{\alpha}\right]. \end{aligned}$$

Because we assumed that \mathcal{D}_s^1 is continuous for each s , we obtain that for each s , by choosing $\alpha \gg \beta$ large enough (ensured by making C large enough), $\text{Prob}_{r \sim \mathcal{D}_s^1}\left[r \leq \frac{\xi}{c_s} - \frac{\beta}{\alpha}\right] \geq \text{Prob}_{r \sim \mathcal{D}_s^1}\left[r \leq \frac{\xi}{c_s}\right] - \epsilon$. By choosing C as the maximum of the corresponding values, this inequality holds simultaneously for all s . In summary, we obtain that

$$\text{Prob}_{(r,d) \sim \mathcal{D}_s^2}[rc_s + d \leq \xi] \geq \text{Prob}_{r \sim \mathcal{D}_s^2}[rc_s \leq \xi] + 2\epsilon.$$

Substituting this inequality into the objective value (14) shows that the objective values (14) and (15) differ by at most

$$\sum_s p_s \cdot (|C(1 - \gamma - B)| \cdot 2\epsilon + \lambda\beta) = |C(1 - \gamma - B)| \cdot 2\epsilon + \lambda\beta.$$

The optimum value of the objective (15) grows at least linearly in C . The reason is that when C is scaled up by any constant ν , a feasible solution is obtained by scaling all c_s up by ν as well; this results in a multiplicative increase of ν in the objective value. Thus, the *best* c_s values must attain at least such an increase. Because the error term $C(1 - \gamma - B)2\epsilon + \lambda\beta$ is at most $O(\epsilon) \cdot \text{OPT}(\mathcal{D}^1)$, we obtain that

$$\text{OPT}(\mathcal{D}^2) \leq \text{OPT}(\mathcal{D}^1) \cdot (1 + O(\epsilon)).$$

Finally, adding in the $O(\epsilon)$ terms for signals not considered in this argument does not change the above conclusion. Making ϵ arbitrarily small (and scaling the Diamonds-in-the-Rough instance correspondingly) then shows in the limit that the distributions over linear and affine functions have the same worst-case behavior.

8 Garbling of Signaling Schemes

In this section, we show that if we garble a signaling scheme φ into φ' , the optimal approximation ratio for the garbled signaling scheme φ' cannot improve. By “garbling,” we mean (stochastically) mapping each signal in φ to a random (new) signal in φ' . Marschak and Miyasawa [26] gave the following¹⁴ formal definition (rephrased to fit our model):

Definition 25 *Let φ, φ' be signaling schemes with respective signal spaces Σ, Σ' . Then, φ' is a garbling of φ if for all conversion ratios r and signals $s \in \Sigma, s' \in \Sigma'$: $f_{s,s'}(r) = f_s(r)$, where $f_s(r)$ is the pdf of conversion ratio r conditioned on signal s .*

As one can see, the garbled signaling scheme φ' contains less information than the original signaling scheme φ . Recall that $1 - p^*(\psi)\gamma$ gives the optimal approximation ratio when the underlying (exogenous) signaling scheme is ψ . Intuitively, more informative signaling schemes should give rise to better approximation ratios, i.e. $1 - p^*(\varphi)\gamma \geq 1 - p^*(\varphi')\gamma$ if φ' is garbled from φ . Theorem 3, restated here, confirms this intuition.

Theorem 3 *Let φ and φ' be two signaling schemes such that φ' is a garbling of φ . Then, $1 - p^*(\varphi)\gamma \geq 1 - p^*(\varphi')\gamma$.*

Proof. Let $p_s, p_{s'}$ be the unconditional probabilities of observing signals s, s' , as defined in Section 2.3, and let $(q_{s'})_{s' \in \Sigma'}$ be the optimal solution for program (5) with signaling scheme φ' . We will show that the program (5) for signaling scheme φ has a feasible solution $(q_s)_{s \in \Sigma}$ such that $\sum_{s \in \Sigma} p_s q_s = \sum_{s' \in \Sigma'} p_{s'} q_{s'}$. This implies that the policy with the garbled signaling scheme cannot outperform the policy with the original signaling scheme.

Let $p_{s,s'}$ be the (joint) probability that signal s is revealed by φ and signal s' is revealed by the garbling φ' . Thus, $p_{s'} = \sum_{s \in \Sigma} p_{s,s'}$. Conditioned on the two revealed signals s, s' , let $q_{s,s'}$ be the probability that the policy with parameters $(q_{s'})_{s' \in \Sigma'}$ (which only observes s') obtains a myopic arm pull. Notice that this probability depends on s : while s does not affect the threshold that is set by the policy, it does affect the distribution of the agent’s conversion ratio, and hence the probability of a myopic pull. Thus, the overall probability of a myopic pull with signal s' can be obtained by summing over all (unobserved) signals s as $p_{s'} q_{s'} = \sum_{s \in \Sigma} p_{s,s'} q_{s,s'}$.

Since $q_{s'}$ is a feasible solution of (5), we have

$$\sum_{s' \in \Sigma'} p_{s'} q_{s'} \geq \lambda \sum_{s' \in \Sigma'} p_{s'} \frac{1 - q_{s'}}{F_{s'}^{-1}(q_{s'})}.$$

Substituting $p_{s'} = \sum_{s \in \Sigma} p_{s,s'}$ and $q_{s'} = \frac{\sum_{s \in \Sigma} p_{s,s'} q_{s,s'}}{p_{s'}}$, we can write

$$p_{s'} \cdot \frac{1 - q_{s'}}{F_{s'}^{-1}(q_{s'})} = \frac{p_{s'}}{F_{s'}^{-1}(q_{s'})} \cdot \left(1 - \frac{\sum_{s \in \Sigma} p_{s,s'} q_{s,s'}}{p_{s'}}\right) = \sum_{s \in \Sigma} p_{s,s'} \cdot \frac{1 - q_{s,s'}}{F_{s'}^{-1}(q_{s'})}.$$

Next, we observe that $F_s^{-1}(q_{s,s'}) = F_{s'}^{-1}(q_{s'})$, as follows: because φ' is a garbling of φ , we get that the conditional distributions satisfy $F_s(r) = F_{s,s'}(r)$. $F_{s'}^{-1}(q_{s'}) =: \tau$ is a threshold chosen by the mechanism with knowledge solely of s' , chosen to achieve a probability of myopic play of

¹⁴There are several equivalent definitions, of which we have chosen the one most suitable for our purposes.

exactly $q_{s'}$. $q_{s,s'}$ denotes the probability of myopic play with this threshold τ , with the additional knowledge that the ungarbled signal was s . Thus, we obtain that $q_{s,s'} = F_{s,s'}(\tau) = F_s(\tau)$, by the garbling property. Taking inverses now shows that $F_{s'}^{-1}(q_{s'}) = \tau = F_s^{-1}(q_{s,s'})$.

Now define the target probabilities for the ungarbled signals as follows: $q_s = \sum_{s' \in \Sigma'} \frac{p_{s,s'}}{p_s} \cdot q_{s,s'}$. We can then write

$$\begin{aligned}
\sum_{s \in \Sigma} p_s q_s &= \sum_{s \in \Sigma} \sum_{s' \in \Sigma'} p_{s,s'} q_{s,s'} \\
&= \sum_{s' \in \Sigma'} \sum_{s \in \Sigma} p_{s,s'} q_{s,s'} \\
&\geq \lambda \sum_{s' \in \Sigma'} \sum_{s \in \Sigma} p_{s,s'} \cdot \frac{1 - q_{s,s'}}{F_s^{-1}(q_{s,s'})} \\
&= \lambda \sum_{s \in \Sigma} \sum_{s' \in \Sigma'} p_s \cdot \frac{p_{s,s'}}{p_s} \cdot \frac{1 - q_{s,s'}}{F_s^{-1}(q_{s,s'})} \\
&\stackrel{(*)}{\geq} \lambda \sum_{s \in \Sigma} p_s \cdot \frac{1 - \sum_{s' \in \Sigma'} \frac{p_{s,s'}}{p_s} q_{s,s'}}{F_s^{-1}(\sum_{s' \in \Sigma'} \frac{p_{s,s'}}{p_s} \cdot q_{s,s'})} \\
&= \lambda \sum_{s \in \Sigma} p_s \cdot \frac{1 - q_s}{F_s^{-1}(q_s)}.
\end{aligned}$$

Here, the inequality labeled (*) followed by semi-regularity of F_s . The inequality derived above implies that q_s is a feasible solution of (5) with signaling scheme φ , and it attains at least the same utility for the principal. This completes the proof of the theorem. \blacksquare

9 Conclusions

We showed that the framework recently proposed by Frazier et al. [9] can be generalized to the case when different agents have different and non-linear tradeoffs for money vs. utility derived from arm pulls. While the generalized framework does not result in as clean a characterization of feasible regions as the original work of Frazier et al. [9], it nonetheless holds true that time-expanded versions of Gittins index policies are optimal in the worst case, and that worst-case examples are of the simple “Diamond-in-the-rough” form.

We needed to assume a technical condition called semi-regularity for our results. Whether time-expanded policies are optimal in the absence of this condition is open.

There are many natural ways in which the model of Frazier et al. [9] could be further generalized. Perhaps most intriguingly, in the model, agents only interact with the mechanism once, whereas it would be natural to assume that the same agents return multiple times. A natural model here would be one of the principal and just *one* agent, who has a different (steeper) time discount $\gamma' < \gamma$ than the principal, and must be incentivized to pull arms with more foresight. In this sense, we analyzed the special case $\gamma' = 0$.

This direction appears quite a bit more difficult to analyze. If agents may return more than once, this opens the door for strategic behavior; it seems possible that an agent may choose a particular arm to pull to help or prevent the principal from learning, in turn affecting possible future payments. This makes this model quite a bit more complicated to analyze.

If one were to try and take the model into a direction of more realism, one could consider specifics of some of the applications listed in the introduction, such as that of an online retailer. In that case, one may have to account for the fact that agents who receive a payment (i.e., discount) on a product may alter their perception or rating of the product.

References

- [1] George A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, pages 488–500, 1970.
- [2] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proc. 36th IEEE Symp. on Foundations of Computer Science*, pages 322–331, 1995.
- [3] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32(1):48–77, 2003. ISSN 0097-5397.
- [4] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. Bandits with knapsacks. In *Proc. 54th IEEE Symp. on Foundations of Computer Science*, pages 207–216, 2013. arxiv.org/abs/1305.2545.
- [5] Bruno Basset, Olivier Gossner, Marco Scarsini, and Shmuel Zamir. Positive value of information in games. *International Journal of Game Theory*, 32(1):17–31, 2003.
- [6] Dirk Bergemann and Juuso Välimäki. Learning and strategic pricing. *Econometrica*, 64(5): 1124–1149, 1996.
- [7] Patrick Bolton and Christopher Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.
- [8] Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
- [9] Peter Frazier, David Kempe, Jon Kleinberg, and Robert Kleinberg. Incentivizing exploration. In *Proc. 16th ACM Conf. on Economics and Computation*, pages 5–22, 2014.
- [10] John C. Gittins. *Multi-Armed Bandit Allocation Indices*. John Wiley and Sons, New York, 1989.
- [11] John C. Gittins and David M. Jones. A dynamic allocation index for the sequential design of experiments. In J. Gani, editor, *Progress in Statistics*, pages 241–266, Amsterdam, 1974. North-Holland.
- [12] John C. Gittins, Kevin D. Glazebrook, and Richard Weber. *Multi-Armed Bandit Allocation Indices*. John Wiley and Sons, New York, 2nd edition, 2011.
- [13] Ashish Goel, Sanjeev Khanna, and Brad Null. The ratio index for budgeted learning, with applications. In *Proc. 20th ACM-SIAM Symp. on Discrete Algorithms*, pages 18–27, 2009.

- [14] Sudipto Guha and Kamesh Munagala. Approximation algorithms for budgeted learning problems. In *Proc. 38th ACM Symp. on Theory of Computing*, pages 104–113, 2007.
- [15] Sudipto Guha and Kamesh Munagala. Multi-armed bandits with metric switching costs. In *Proc. 36th Intl. Colloq. on Automata, Languages and Programming*, pages 496–507. Springer, 2009.
- [16] Sudipto Guha and Kamesh Munagala. Approximation algorithms for bayesian multi-armed bandit problems. arxiv.org/abs/1306.3525, 2013.
- [17] Jack Hirshleifer. The private and social value of information and the reward to inventive activity. *The American Economic Review*, pages 561–574, 1971.
- [18] Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. Adaptive contract design for crowdsourcing markets: Bandit algorithms for repeated principal-agent problems. In *Proc. 16th ACM Conf. on Economics and Computation*, pages 359–376, 2014.
- [19] Michael N. Katehakis and Arthur F. Veinott, Jr. The multi-armed bandit problem: decomposition and computation. *Math. Oper. Res.*, 12(2):262–268, 1987. ISSN 0364-765X.
- [20] Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Multi-armed bandits in metric spaces. In *Proc. 39th ACM Symp. on Theory of Computing*, pages 681–690, 2008.
- [21] Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the “wisdom of the crowd”. In *Proc. 15th ACM Conf. on Electronic Commerce*, pages 605–606, 2013.
- [22] Tze Leung Lai and Herbert E. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [23] Ehud Lehrer, Dinah Rosenberg, and Eran Shmaya. Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2):670–682, 2010.
- [24] David G. Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1968.
- [25] Yishay Mansour, Aleksandrs Slivkins, and Vasilis Syrgkanis. Bayesian incentive-compatible bandit exploration. In *Proc. 17th ACM Conf. on Economics and Computation*, pages 565–582, 2015.
- [26] Jacob Marschak and Koichi Miyasawa. Economic comparability of information systems. *International Economic Review*, 9(2):137–174, 1968.
- [27] Hebert E. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58:527–535, 1952.
- [28] Adish Singla and Andreas Krause. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *22nd Intl. World Wide Web Conference*, pages 1167–1178, 2013.
- [29] Aleksandrs Slivkins and Jennifer Wortman Vaughan. Online decision making in crowdsourcing markets: Theoretical challenges (position paper). *ACM SIGecom Exchanges*, 12(2), 2013.
- [30] Michael Spence. Job market signaling. *The Quarterly Journal of Economics*, pages 355–374, 1973.

- [31] George J. Stigler. Information in the labor market. *The Journal of Political Economy*, pages 94–105, 1962.
- [32] Peter Whittle. Multi-armed bandits and the Gittins index. *Journal of the Royal Statistical Society. Series B (Methodological)*, 42(2):143–149, 1980.

A Missing Proofs

Lemma 7 *Let G be the CDF of a non-negative random variable, and define $G^{-1}(x) := \sup\{t \geq 0 : G(t) \leq x\}$. If $G^{-1}(x)(1-x)$ is concave, then $\frac{1-x}{G^{-1}(x)}$ is convex. In particular, regularity implies semi-regularity.*

Proof. We only need to look at 3 points x , $\alpha x + (1-\alpha)y$ and y , where $0 \leq x < y \leq 1$. Let $z = \alpha x + (1-\alpha)y$. By concavity of $G^{-1}(x)(1-x)$, we have

$$G^{-1}(z)(1-z) \geq \alpha G^{-1}(x)(1-x) + (1-\alpha)G^{-1}(y)(1-y) \quad (16)$$

Rearranging, we have:

$$0 \geq \alpha(1-x)\frac{G^{-1}(x) - G^{-1}(z)}{G^{-1}(z)} + (1-\alpha)(1-y)\frac{G^{-1}(y) - G^{-1}(z)}{G^{-1}(z)} \quad (17)$$

By monotonicity of G^{-1} , $G^{-1}(x) \leq G^{-1}(z) \leq G^{-1}(y)$. This implies

$$0 \geq \alpha(1-x)\frac{G^{-1}(x) - G^{-1}(z)}{G^{-1}(x)} + (1-\alpha)(1-y)\frac{G^{-1}(y) - G^{-1}(z)}{G^{-1}(y)}, \quad (18)$$

as $G^{-1}(x) \leq G^{-1}(z)$ and $G^{-1}(y) \geq G^{-1}(z)$. One can see that inequality (18) is equivalent to the convexity of $\frac{1-x}{G^{-1}(x)}$. ■

The proofs are very straightforward (syntactic) modifications of those of the corresponding lemmas in [9], we include them here for completeness.

Lemma 11 (Modification of Lemma 4.2 of [9]) *Given a parameter λ and a signaling scheme φ . Let $\zeta_{t-1} = \sum_{t' < t} Z_{t'}$ be the total number of non-myopic steps performed by the time-expanded algorithm $\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}$ prior to time t , where \mathbf{q} satisfies $\sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)}$. Then, for any $0 \leq n \leq t$,*

$$\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \zeta_{t-1} = n] \geq \mathbb{E}_{\mathcal{A}} [v_{n, i_n}].$$

Proof. The Lagrangian utility at time t is $v_{t, i_t} - \lambda c_{t, i_t}$. Let $x = \max_i \mathbb{E} [v_{t, i} \mid \mathbf{S}_t, Z_{0:t-1}]$ and $y = \mathbb{E} [v_{t, i_t}^* \mid \mathbf{S}_t, Z_{0:t-1}]$, where $Z_{0:t-1} = (Z_0, Z_1, \dots, Z_{t-1})$. Since the myopic arm is played with probability q_s , the expected Lagrangian utility will be

$$\begin{aligned} \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \mathbf{S}_t, Z_{0:t-1}] &= \sum_{s \in \Sigma} p_s (q_s x + (1-q_s)(y - \lambda \frac{x-y}{F_s^{-1}(q_s)})) \\ &= y \sum_{s \in \Sigma} p_s + (x-y) \sum_{s \in \Sigma} p_s (q_s - \frac{\lambda(1-q_s)}{F_s^{-1}(q_s)}) \geq y. \end{aligned}$$

The last inequality above is due to $x \geq y$ (by the myopic arm preference) and $\sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)}$ (by assumption of the lemma). Notice the condition $\sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)}$ in the statement of the lemma allowed us to cancel out the myopic arm-rewards x , which would otherwise have been difficult to analyze. Now we have

$$\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \mathbf{S}_t, Z_{0:t-1}] \geq \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t^*} \mid \mathbf{S}_t, Z_{0:t-1}].$$

Since the expected reward only depends on the current state of arm \mathbf{S}_t , we can add the status of previous round to the condition,

$$\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \mathbf{S}_{0:t}, Z_{0:t-1}] \geq \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t^*} \mid \mathbf{S}_{0:t}, Z_{0:t-1}].$$

Taking conditional expectation on both sides with the condition $\hat{\mathbf{S}}_t, \zeta_{t-1} = n$,

$$\begin{aligned} & \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} \left[\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \mathbf{S}_{0:t}, Z_{0:t-1}] \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n \right] \\ & \geq \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} \left[\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t^*} \mid \mathbf{S}_{0:t}, Z_{0:t-1}] \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n \right]. \end{aligned}$$

Since $\hat{\mathbf{S}}_t, \zeta_{t-1} = n$ is measurable by $\mathbf{S}_{0:t}, Z_{0:t-1}$, we can use the tower property of conditional expectations and get

$$\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n] \geq \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t^*} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n].$$

Let $\tau \geq t$ be the next time when $\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}$ performs a non-myopic pull. Since the reward sequence for each arm forms a Martingale, and i_u^* is fixed with a certain $\hat{\mathbf{S}}_t$, we have that $\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{u, i_u^*} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n]$ is identical for all $u \geq t$. Writing $\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{\tau, i_\tau^*} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n]$ with the law of total probability, we have

$$\begin{aligned} \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{\tau, i_\tau^*} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n] &= \sum_{u \geq t} \text{Prob}[\tau = u] \cdot \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{u, i_u^*} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n, \tau = u] \\ &= \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t^*} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n]. \end{aligned}$$

So far, we have proved that

$$\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n] \geq \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{\tau, i_\tau^*} \mid \hat{\mathbf{S}}_t, \zeta_{t-1} = n].$$

Taking conditional expectations with respect to $\zeta_{t-1} = n$, then applying the law of iterated expectation,

$$\mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \zeta_{t-1} = n] \geq \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{\tau, i_\tau^*} \mid \zeta_{t-1} = n].$$

Note the right-hand side is exactly $\mathbb{E}_{\mathcal{A}} [v_{n, i_n}]$ by definition of $\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}$, finishing the proof. \blacksquare

Lemma 12 (Variation of Lemma 3.2 of [9]) *Given a parameter λ and a signaling scheme φ . Assume \mathbf{q} satisfies $\sum_{s \in \Sigma} p_s q_s \geq \lambda \sum_{s \in \Sigma} p_s \frac{1-q_s}{F_s^{-1}(q_s)}$, then for $\eta = \frac{(1-p)\gamma}{1-p\gamma}$, where $p = \sum_{s \in \Sigma} p_s q_s$, we have*

$$R_\lambda^{(\gamma)}(\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}) \geq \frac{1-\eta}{1-\gamma} \cdot R^{(\eta)}(\mathcal{A}).$$

Proof.

$$\begin{aligned}
R_{\lambda}^{(\gamma)}(\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}) &= \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t}] \\
&= \sum_{t=0}^{\infty} \sum_{n=0}^{\infty} \gamma^t \mathbb{E}_{\mathbf{TES}_{\mathbf{q}, \mathcal{A}, \varphi}} [v_{t, i_t} - \lambda c_{t, i_t} \mid \zeta_{t-1} = n] \cdot \text{Prob}[\zeta_{t-1} = n] \\
&\stackrel{\text{Lemma 11}}{\geq} \sum_{n=0}^{\infty} \mathbb{E}_{\mathcal{A}} [v_{n, i_n}] \cdot \sum_{t=0}^{\infty} \gamma^t \cdot \text{Prob}[\zeta_{t-1} = n] \\
&= \sum_{n=0}^{\infty} \mathbb{E}_{\mathcal{A}} [v_{n, i_n}] \cdot \sum_{t=0}^{\infty} \gamma^t \cdot \binom{t}{n} \cdot (1-p)^n p^{t-n} \\
&= \sum_{n=0}^{\infty} \mathbb{E}_{\mathcal{A}} [v_{n, i_n}] \cdot \gamma^n (1-p)^n \cdot \sum_{i=0}^{\infty} \binom{n+i}{n} \cdot (\gamma p)^i \\
&= \sum_{n=0}^{\infty} \mathbb{E}_{\mathcal{A}} [v_{n, i_n}] \cdot \gamma^n (1-p)^n \cdot (1-\gamma p)^{-(n+1)} \\
&= \sum_{n=0}^{\infty} \mathbb{E}_{\mathcal{A}} [v_{n, i_n}] \cdot \frac{1-\eta}{1-\gamma} \cdot \eta^n.
\end{aligned}$$

In the penultimate step, we use $\sum_{i=0}^{\infty} \binom{n+i}{n} x^i = (1-x)^{-(n+1)}$ for $n \geq 0$ and $|x| < 1$. ■

B Impossibility of Exploration with Reward-Dependent Payments

In this section, we show that if the payment scheme is a function only of the arm reward obtained by the agent, it can be impossible to incentivize the agent to pull the optimal arm.

We use a slightly modified “Diamonds in the Rough” instance $\Delta(B, \gamma)$, by changing the myopic arm to two arms. Both of them produce i.i.d. rewards from a known distribution. The distributions are showed in Table 1:

Reward	Arm 1 outputs reward with probability	Arm 2 outputs reward with probability
$(1-\gamma)B \cdot M$	$(1+\epsilon)/M$	$(1-\epsilon)/M$
0	$1 - (1+\epsilon)/M$	$1 - (1-\epsilon)/M$

Table 1: Distributions of myopic arms

We emphasize that this is not the type of arm with degenerate distribution that produces an initially unknown constant reward.

Also we keep the infinite supply of arms with degenerate distribution. Recall the definition as follows,

1. With probability $1/M$, the arm’s reward is a degenerate distribution of the constant $(1-\gamma)B \cdot M$ (good state);
2. With probability $1 - 1/M$, the arm’s reward is a degenerate distribution of the constant 0 (bad state).

Because the payment function can depend only on the reward that the agent obtained, it is completely characterized by the payment p_H in response to obtaining $(1-\gamma)B \cdot M$ and the payment p_L in response to obtaining 0. When $(1-\gamma)B \cdot M + p_H \geq p_L$, we have

$$((1-\gamma)B \cdot M + p_H) \cdot \frac{1+\epsilon}{M} + p_L \cdot \left(1 - \frac{1+\epsilon}{M}\right) \geq ((1-\gamma)B \cdot M + p_H) \cdot \frac{1}{M} + p_L \cdot \left(1 - \frac{1}{M}\right);$$

otherwise, we have

$$((1-\gamma)B \cdot M + p_H) \cdot \frac{1-\epsilon}{M} + p_L \cdot \left(1 - \frac{1-\epsilon}{M}\right) \geq ((1-\gamma)B \cdot M + p_H) \cdot \frac{1}{M} + p_L \cdot \left(1 - \frac{1}{M}\right).$$

So the expected utility of pulling the optimal arm is always weakly dominated by one of two myopic arms. Therefore, there is no way to incentivize the agent to pull the optimal arm.